Remarks

Favorable reconsideration of this application is requested in view of the following remarks. For the reasons set forth below, Applicant respectfully submits that the claimed invention is allowable over the cited references.

The Advisory Action dated December 13, 2004 indicated that the previously proposed amendments are not entered, claim 23 is or will be allowable, and that claims 1-22 are rejected.

After the cancellation of claim 19, claims 1-9, 13, 14 and 20 are rejected under 35 U.S.C. § 103(a) over *Edholm* (U.S. Patent No. 6,449,269) in view of teaching in the art; claim 10 is rejected under 35 U.S.C. § 103(a) over *Edholm* and teaching in the art and further in view of *Bertin* (U.S. Patent No. 6,097,243); claim 11 is rejected under 35 U.S.C. § 103(a) over *Edholm* and teaching in the art and further in view of *Mason* (U.S. Patent No. 6,272,451); claims 12 and 15-18 are rejected under 35 U.S.C. § 103(a) over *Edholm* and teaching in the art and further in view of *Maeda* (U.S. Patent No. 5,884,074); and claims 21-22 are rejected under 35 U.S.C. § 103(a) over *Edholm* and teaching in the art and further in view of *Adelman* (U.S. Patent No. 5,598,362).

Applicant appreciates the Examiner's indication that claim 23 is or will be allowable. In this regard, claim 23 has been amended to remove its dependency and to include all the limitations in the base and intervening claims from which it depends. Applicant understands that claim 23, as amended, is allowable.

New claims 24 – 28 have been added. The subject matter of new claims 24 – 28 does not introduce new matter. The claims are believed to be patentable over the cited references, which do not teach or suggest the claimed invention as so characterized.

Claims 1, 2, 4 - 6, 8, 9, 14 - 17, and 20 have been amended to more clearly claim the invention. In addition, claim 13 has been amended in part to more clearly claim the invention. These amendments have not been made in view of any cited art and are not intended to narrow the scope of the claims.

Referring to claim 1 as an example, the claimed programmable audio processor chip comprises:

a DSP voice compression device adapted to compress the voice data;

audio processing circuitry programmed with an audio processing software application for processing the compressed voice data;

an IP network stack adapted to store and process IP data, the IP network stack data including protocols for processing the compressed voice data via an IP network; and

a communication stack adapted to store and process communications data, the communications stack data including audio processing protocols for processing the compressed voice data.

Applicant respectfully disagrees that the '269 reference teaches or suggests the limitations directed to a programmable audio processor chip for processing voice data that includes an IP network stack and/or a communication stack. Contrary to the interpretation in the final Office Action dated July 23, 2004, the '269 reference does not teach a chip with either an IP network stack or a communication stack, which each include a protocol layer higher than the IP layer; rather, the '269 reference teaches an ASIC having the IP layer as its highest layer (see col. 2, line 44-46 and col. 3, line 5-17 of the '269 reference). The final Office Action dated July 23, 2004 states on page 13: "Moreover, Edholm'269 discloses wherein IP network stack includes at least one of: a TCP/IP stack and a H.323 stack (see col. 11, line 11-26, and col. 12, line 1-12)". Applicant agrees with Examiner's assessment that an IP network stack includes a protocol layer higher than the layer 3 protocol of the IP layer, such as the layer 4 protocol of TCP; however, col. 11, line 11-26 and col. 12, line 1-12 of the '269 reference clearly state that higher OSI layer (layer 4+) signaling (e.g. TCP/IP and/or ITU H.323) are processed by a separate phone server operating in tandem with the IP telephone. The ASIC claimed in the '269 reference claim 4, which is directed to certain portions of the IP telephone (specifically, "said finite state machine, said memory, said packetizer, and said network interface") and not the separate phone server, clearly does not include any layers higher than the IP layer, and thus does not include an IP network stack, which includes a protocol layer higher than the IP layer. Further, because the communication stack also includes a protocol layer higher than the IP layer, the '269 reference also does not teach a chip including a communication stack. For further information regarding this terminology, reference may be made to "The Irwin Handbook of Telecommunications" (1997) and U.S. Patent 6,651,117, "Network Stack Layer Interface", the former attached hereto.

Independent claims 1, 20, 23, and 25 all recite a programmable audio processor chip including an IP network stack and a communication stack. The remaining independent claim 13 is amended to also recite a programmable audio processor chip including an IP network stack and a communication stack. The IP network stack and the communication stack include one or more layers higher than the IP layer, and are thus all independent claims are distinguished from the '269 reference.

In view of the above discussion, Applicant believes that the rejections have been overcome and the application is in condition for allowance. A favorable response is requested. Should there be any remaining issues that could be readily addressed over the telephone, the Examiner is encouraged to contact the undersigned at (651) 686-6633.

Respectfully submitted,

CRAWFORD MAUNU PLLC
1270 Northland Drive, Suite 390
St. Paul, MN 55120
651/686-6633

Dated: January 12th, 2005

By: __________________

Robert J. Crawford
Reg. No. 32,122
Eric J. Curtin
Reg. No. 47,511

Attachment: The Irwin Handbook of Telecommunications, p.p. i-iii, 98-109

Third Edition

THE IRWIN HANDBOOK OF TELECOMMUNICATIONS

JAMES HARRY GREEN

IRWIN
Professional Publishing®
Chicago • London • Singapore

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold with the understanding that neither the author nor the publisher is engaged in rendering legal, accounting, or other professional service. If legal advice or other expert assistance is required, the services of a competent professional person should be sought.

From a Declaration of Principles jointly adopted by a Committee of the American Bar Association and a Committee of Publishers.

Times Mirror
Higher Education Group

Library of Congress Cataloging-in-Publication Data

Green, James H. (James Harry)
The Irwin handbook of telecommunications / James Harry Green.
p. cm.
Includes index.
ISBN 0-7863-0479-0
1. Rev. ed. of: The Irwin handbook of telecommunications management. 1995. 2. Telecommunication systems—United States. 3. Business enterprises—United States—Communication systems. I. Green, James H. (James Harry) Irwin handbook of telecommunications management. II. Title.
TK5102.3.U6G74 1997 96–23791
621.382—dc20

Printed in the United States of America
1 2 3 4 5 6 7 8 9 0 BB 3 2 1 0 9 8 7 6

PREFACE

These are exciting times for people in the telecommunications business. The breakup of the Bell System in 1984 trigged an onslaught of change greater than the sum of all change from the invention of the telephone more than 100 years earlier, and still the change continues. AT&T has broken itself into three more companies, voluntarily this time. Cisco purchased Stratacom, and IBM acquired Lotus. The telecommunications and media giants dance with one another looking for combinations to help them achieve critical mass. Time-Warner, for example, found a successful fusion with Turner Broadcasting, but not with US West. Congress has passed new telecommunications legislation that will cost local exchange companies their monopolies, while opening the way for them to enter new markets, and that has spawned more changes such as the pairs of ex-Bell companies that are combining. Southwestern Bell Corp. has merged with Pacific Bell, and Hynex and Bell Atlantic have announced plans to combine.

This has been a wrenching time for those accustomed to the old ways of doing business: guaranteed returns on investment, rigid and protective tariffs, and government regulation in lieu of competition. Hundreds of thousands of people have found their once steady and predictable jobs realigned or lost in the tumult of change. Service quality that once was taken for granted has often been all but forgotten in the jockeying for competitive position.

But a revolution without discomfort is no revolution at all, and what we are witnessing is revolutionary change in an industry that once was accustomed to stability. Hundreds of new companies and countless products have emerged since the Bell System breakup, and we have barely seen the beginnings of what human ingenuity can devise when the shackles are removed.

Since the second edition of this book was published, many changes have occurred that affect telecommunications. Then, the very term asynchronous transfer mode (ATM) was unknown to most within the industry. Today it has become a mantra. Although ATM has still not been widely deployed, nor are standards complete, it has generated more interest than any technology in recent memory. Structured wiring standards had not been developed, computer-telephony integration was unknown to most users, and LAN internets were in their infancy. Today these technologies are either mainstream or rapidly gaining acceptance.

Multi-Tech Systems, Inc.
Paradyne Corp.
RAD Data Communications

Video Display Terminals

Digital Equipment Corp.
Harris Corp.
IBM Corp.
Wyse Technology
Unisys Corp.

CHAPTER 4 DATA COMMUNICATIONS PROTOCOLS

When diplomats speak of protocols they refer to the etiquette, customs, and procedures by which political relations are conducted. In the telecommunications world protocols fill a similar purpose. Devices exchange signals, sometimes called a "handshake," that establish the terms and conditions under which they will communicate. Unlike human beings who have the ability to adjust to an unfamiliar protocol, however, data devices are unable to communicate at all unless their protocols match within narrowly defined limits.

Protocol compatibility and standardization are the most important issues in data communications. Most major computer manufacturers have developed proprietary protocols that are incompatible with those of other manufacturers. Even standard protocols developed by international agencies, such as ITU's HDLC and X.25, provide multiple options and are not always interchangeable between applications.

Incompatible protocols can communicate with each other through *protocol converters*. A protocol converter known as a *gateway* communicates with both connecting protocols using their languages. Certain value-added networks also offer protocol conversion. They communicate with the DTE using that DTE's own protocol, convert it to the network protocol, and transport it to the distant DTE in its own language. The incompatibility of protocols has long been a stumbling block in the path of full interconnectability of data networks. Although international standards have progressed significantly, the problems of incompatibility are likely to remain for many years.

The simplest protocols have been established by common usage. For example, asynchronous protocol is used by every device that has a standard serial interface. It is universally accepted, and when the device is initialized with the proper speed, parity, stop bits, and data bits, it can communicate with any other device that is identically set up. Such simplicity is achieved at a price, however. As we have seen, the asynchronous protocol does not provide for such an important function as error correction. The lack of error correction was a serious drawback of the asynchronous protocol until development of file transfer protocols such as X-Modem, Kermit, and the like. Now, many modems have built-in error correction using another protocol, Microcom's MNP, which is the international standard V.42.

The hundreds of different protocols in the data communications industry are at the root of traits that make data communications more complex than voice. Not only is the sheer number of protocols enormous, but most of them have options and variations that must be satisfied before devices can communicate. Many protocols are proprietary. The major computer manufacturers have proprietary networks derived from protocols of their design. Most have been released into the public domain or licensed to allow other manufacturers to interoperate, but the standard remains under the manufacturer's control.

This chapter discusses the kinds of functions data protocols perform and the terminology that is common to most protocols. We will discuss the International Standards Organization's Open Systems Interconnect model (OSI), which is at the heart of most current protocol designs. We will review two examples of protocols that are in common use today: transmission control protocol/internet protocol (TCP/IP) and point-to-point protocol (PPP). These protocols are complex, and will be reviewed in concept only as a way of illustrating how protocol functions are carried out. For a more detailed bit-level discussion of these and other protocols, readers are referred to publications listed in the bibliography. Subsequent chapters will discuss many other protocols in the context in which they are used.

PROTOCOL TERMINOLOGY AND FUNCTIONS

Data protocols may be implemented in firmware (a chip), hardware, or a combination of both. Just as computer programs are usually written in modules to simplify administration, layered protocols allow developers to write software to a clearly defined interface. Each layer has a defined function. If a function or specification changes, it isn't necessary to change the entire protocol stack; only the affected layer and its interfaces with other layers are changed.

A good example of layered protocols in action is found in LAN standards. These standards, which are discussed in more detail in Chapter 29, are designed to work at the first two layers of the OSI model. Hardware vendors can build network interface cards (NICs) to connect to any of the transmission media that LAN standards support (twisted-pair cable, fiber optics, coaxial cable, and so on). Network operating system developers such as Novell, Banyan, Microsoft, and others can develop higher layer protocols that talk to any NIC. The card manufacturers provide software drivers to enable the functions in their cards to communicate with the network operating system.

LAN standards further illustrate how the protocol is deployed. The portion of the protocol that describes how the NIC communicates with the transmission medium is implemented in firmware. The card manufacturers use chips that implement the access protocol. The drivers and all of the network operating system are implemented in software. The result is a complete network that can be composed of NICs from one manufacturer, bridges or routers from another, a network operating system from a third manufacturer, and computers from a mixture of companies. Although interoperability is not completely assured, manufacturers are able to design to known interfaces, and when problems occur the standard makes it easier to determine what corrective action is appropriate.

Protocol Functions

To continue with the analogy of a diplomatic protocol, the way people behave in a situation is dictated by a complex set of rules. Diplomatic protocols suggest who is seated next to whom, how officials of different ranks are to be addressed, what kind of response is appropriate to another's statement, who is introduced to whom, and other such niceties that govern diplomatic affairs. Data protocols dictate some of the same types of relationships. In this section we will discuss some of the major functions of protocols. In a layered protocol these functions are usually assigned to one layer, but the rules regarding this are not rigid. For example, every Ethernet NIC has an address embedded in firmware. That address permanently belongs to the station, and is not duplicated anywhere else in the world. That address is used to identify the station in intraLAN transactions. If the LAN is connected to a distant LAN, the internetworking protocol is likely to be TCP/IP. IP, a higher level protocol, uses an addressing scheme entirely different from that used on the LAN. We will discuss TCP/IP later in this chapter. In this section we will discuss the major functions of protocols without concern for where or how they are implemented.

Session Control

The major objective of interactions between protocols is to establish a session across a network. A session begins when devices establish communication, and ends when the communication terminates. In dial-up communications a session begins when parties establish a connection across the network, and ends when one party hangs up. In data communications a session may begin when a user logs on a distant computer, and ends with log-off. The protocol authenticates the parties before permitting communication to begin.

Data networks handle sessions in two distinct ways: *connectionless* and *connection-oriented*. In a connection-oriented protocol the devices have a physical or logical connection across the network; the connection is set up at the start of the session and remains for its duration. The connection can be circuit-switched as in the telephone network, or it can be a *virtual* connection, which is defined in a software path that shares circuits with other sessions.

A connectionless session is one in which data is launched into the network and delivered to the distant end based on its address. The postal service is an example of a connectionless operation. The user doesn't care how a letter gets to its destination. Each letter is individually addressed and handed to the post office for delivery to the addressee. Most LANs are also connectionless. A stream of information is launched onto the network. All stations copy the message but retain it only if they are the addressee. In a data network, connectionless operation means that each packet or frame must contain the address of the sending and receiving stations. In a connection-oriented session the packets or frames typically contain a path identifier, but do not need the address of either the sender or the receiver after the session is set up.

Communications Control

Protocols can be classified as peer-to-peer or master-slave. In the latter protocol the master controls the functioning of the data link and controls data transfer between the host and its terminals. All communication between slaves goes through the master. A peer-to-peer protocol does not use a controller, so devices can communicate with one another at will.

Link Management

After the session is set up, the protocol controls the flow of data across the data link.

Synchronizing

At the start of a session data devices exchange signals to determine such variables as the bit rate of the modems, whether compression or error correction will be used, and so on. Modems exchange signals to determine the highest speed at which they can exchange data, falling back to a lower speed if the circuit will not support the maximum.

Addressing

Every session requires an address to set up a connection if the protocol is connection-oriented, or to route packets if it is connectionless. Not all protocols contain addresses. Many of them rely on higher or lower layers for addressing.

Routing

In data networks having multiple routes to the destination, the protocol determines the appropriate route based on variables such as cost, congestion, distance, and type of facility.

Data Segmenting and Reassembly

A continuous data stream from the source is segmented into frames, cells, or packets as appropriate and equipped with header and trailer records for transmission over the network. At the distant end the protocol strips the overhead bytes and reassembles the data stream for delivery to the receiver.

Data Formatting

The bit stream may require conditioning before transmission and restoration after reception. For example, conditioning could include encryption or compression.

Flow Control

Protocols protect networks from congestion by sending signals to the source to halt or limit traffic flow.

Supervision

The protocol establishes a connection, determines how the session will be started and ended, which end will control termination of the session, how charging will be handled, and so on.

Error Detection and Correction

Protocols check for errors, acknowledge correctly received data blocks, and send repeat requests when blocks contain an error. In the most primitive type of error correction, each packet is separately acknowledged, so the sending device must wait for an acknowledgment before sending another packet. More sophisticated protocols can acknowledge multiple packets using one of two types of acknowledgment. A *selective repeat* acknowledgment enables the receiving device to request specific packets to be repeated. In the *go-back-n* method the receiver instructs the sender to resend an errored packet and all subsequent packets.

Failure Recovery

If the session terminates unexpectedly, the protocol determines how to keep the application from being corrupted.

Sequencing

If data blocks are received out of their original sequence, the protocol delivers them to the receiving device in the correct order.

Setting Session Variables

The protocol determines such variables as whether the session will be half or full duplex, network login and authentication, file transfer protocols that will be used, and so on.

THE OPEN SYSTEMS INTERCONNECT MODEL

The standardization efforts in protocols have resulted in layered protocols. A layer is a discrete set of functions the protocol is designed to accomplish. The International Standards Organization (ISO) has published a seven-layer protocol model, the Open Systems Interconnect (OSI) model, which is illustrated in Figure 4–1.

Controlling communications in layers adds some extra overhead because each layer communicates with its counterpart through header records, but layered protocols are easier to administer than single-layer protocols and provide greater opportunity for standardization. Although protocols are complex, functions in each layer can be modularized so the complexity can be dealt with separately by system designers. Layered control offers an opportunity for standardization and interconnection between the proprietary architectures of different manufacturers. Generally, the degree of standardization is greatest at the first layer, and becomes increasingly disparate in the higher layers. The seven OSI layers are defined below. Table 4–1 lists the OSI layers and some of the standards that apply to each layer. The higher layers in the protocol stack are more abstract and in some cases less well-defined than the first three or four layers, which have been in common use for many years.

FIGURE 4–1

International Standards Organization Open Systems Interconnection Model

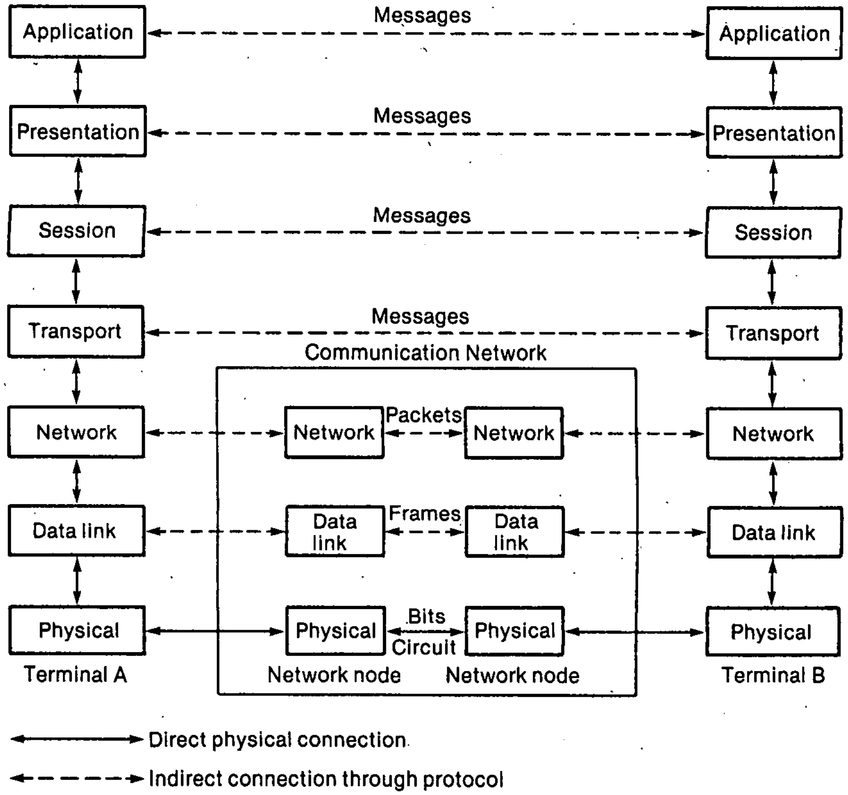

Layer 1—Physical

The first layer describes the method of physical interconnection over a circuit. The physical layer contains the rules for the transmission of *bits* between machines and standardizes pin connections between DCE and DTE. The standards discuss modulation methods and multiplexing over the physical medium, which is wire, fiber optics, coaxin cable, or wireless. For example, EIA-232 is a common standard for serial port connections. Its speed and distance limitations are overcome by using balanced interfaces such as EIA-422 or V.35. Two devices can communicate using nothing but the physical layer. For example, if the serial ports of two computers are connected through an adapter

TABLE 4–1

Representative Protocols of the OSI Layers

Layer	Common Standards
1. Physical	EIA-232 EIA-422 V.35
2. Data Link	High-level data link control (HDLC) Balanced link access procedure (LAPB) Designated link access procedure (LAPD) IEEE 802.2 logical link control
3. Network	X.25 packet level protocol Internet protocol (IP) Connectionless network protocol (CLNP) Address resolution protocol (ARP) IBM/SNA Path control
4. Transport	Transmission control protocol (TCP) User datagram protocol (UDP) Netware control protocol (NCP) ISO transport protocol
5. Session	ISO connection-oriented session protocol NETBIOS IBM/SNA data flow control
6. Presentation	ISO connection-oriented presentation protocol Microsoft server message block protocol Netware file service protocol
7. Application	X.400 Message handling service (MHS) ISO file transfer, access, and management (FTAM) ISO office document architecture (ODA) ISO virtual terminal service Simple mail transfer protocol (SMTP) Virtual terminal (TELNET)

known as a *null modem,* they can send data to each other. The null modem connects the transmitting data and signaling leads of each computer to the corresponding receiving leads of the other.

Layer 2—Data Link

Data link protocols are concerned with the transmission of *frames* of data between devices. The protocol in the data link layer detects and corrects errors so the user gets an error-free circuit. The data link layer takes raw data characters, creates frames of data from them, and processes acknowledgment messages from the receiver. When frames are lost or mutilated, the logic in this layer arranges retransmission.

Protocols contain flags and headers so DTE can recognize the start and end of a frame. A frame of information, as Figure 3–6 shows, has flags to signal the beginning and ending of the frame, a header containing address and control information, an information field, and a trailer containing CRC bits for error correction. The principal international standard, high-level data link control, has numerous subsets, of which balanced-link access procedure and designated-link access procedure are common. The former is used in packet-switched data networks, and the latter as the access protocol for ISDN.

Layer 3—Network

The network layer accepts messages from the higher layers, breaks them into *packets,* routes them to the distant end through the link and physical layers, and reassembles them in the same form in which the sending end delivered them to the network. The network layer controls the flow of packets, controls congestion in the network, and routes between nodes to the destination. The X.25 protocol is a connection-oriented protocol for access to a packet-switched data network. Internet protocol (IP) is one of the most widely used protocols in the world which, together with transmission control protocol, forms a common language used in most internets.

Layer 4—Transport

The transport layer controls end-to-end integrity between DTE devices, establishing and terminating the connection. It segments data into manageable protocol data units (PDUs), and reassembles them at the receiving end. It is responsible for flow control and end-to-end error correction. If the lower layers have any shortcomings in their ability to deliver data with complete integrity, it falls to the transport layer to overcome them. Transmission control protocol, (TCP), which is discussed later in this chapter, is the most widely used transport layer protocol. User datagram protocol (UDP) is a connectionless protocol that is used by simple network management protocol (SNMP).

Layer 5—Session

The user communicates directly with the session layer, furnishing an address that the session layer converts to the address the transport layer requires. The conventions for the session are established in this layer. For example, the validity of the parties can be verified by passwords. The session can be established as a full-duplex or a half-duplex session. The session layer determines whether machines can interrupt one another. It establishes how to begin and terminate a session, and how to restore or terminate the connection in case a failure interrupts the session. If a user attempts a file transfer, for example, the file must be opened, the data moved across the network, and the file closed at the end of the session. If anything happens to disrupt the transfer, the file could be left in limbo. It is the job of the session layer to ensure that the transfer is as orderly as if the distant device was directly connected to the host.

Layer 6—Presentation

This layer interprets the character stream that flows between terminals during the session. For example, if encryption or bit compression is used, the presentation layer may provide it. This layer is the least well developed of the OSI model, and is skipped in many implementations.

Layer 7—Application

The application layer is the interface between the network and the application running on the computer. Examples of application layer functions now in use are ITU's X.400 electronic mail protocol and its companion X.500 directory services protocol. Message-handling service (MHS) is an important protocol for enabling X.400 E-mail systems to communicate. ISO's file transfer, access, and management (FTAM) is a protocol for managing and manipulating files across a network. Other protocols include virtual terminal (VT), which provides a standard terminal interface, and electronic document interchange (EDI), which uses the MHS platform for transferring electronic documents across networks.

The objective of the OSI reference model is to establish a framework that will allow any conforming system or network to connect and exchange signals, messages, packets, and addresses. The model makes it possible for communications to become independent of the manufacturer that devised the technology, and to shield the user from the need to understand the complexity of the network. It should be understood that although the OSI model can be used to develop standards, it is not a standard itself. Manufacturers are increasingly announcing support for OSI, and proprietary networks may eventually evolve into compatible standard networks.

TRANSMISSION CONTROL PROTOCOL/INTERNET PROTOCOL (TCP/IP)

TCP/IP is a collection of protocols that were developed in the late 1970s by the Department of Defense as a way of providing interoperability among equipment manufacturers. The protocols emerged from research that

spanned three decades under the auspices of Advanced Research Projects Agency (ARPA). The agency is now called Defense Advanced Research Projects Agency (DARPA). ARPANET, as it was called in its early days, was a loosely confederated collection of networks operated by colleges, universities, and defense-related companies and agencies. Unlike OSI, TCP/IP is not a true international standard, although it is an open standard that is widely used internationally. The standard is administered through the Internet Engineering Task Force, which is a voluntary body. The IETF distributes its recommendations through Internet Requests for Comments, which are open to anyone.

In late 1989 the original ARPANET gave way to a network that now is known as the *Internet.* Internet is a collection of independent packet-switched networks that are interconnected to act as a coordinated unit. Governmental agencies, military branches, educational institutions, and commercial companies operate the networks, but no single body has overall control. The protocols compensate for the unreliability of the underlying networks and insulate users from the need to understand the network's architecture and addressing scheme. Internet has four primary purposes:

- To provide electronic mail service to the users.
- To support file transfer between hosts.
- To permit users to log on to remote computers.
- To provide users with access to information databases.

A TCP/IP internet fits in a three-layer framework atop the physical and data link layers, as shown in Figure 4–2. The application services layer defines the interface to the basic network, which consists of the transmission control and internet protocols. The transport layer has two separate protocols, TCP, which is a connection-oriented protocol, and UDP, which is connectionless. Figure 4–2 compares TCP/IP to the OSI model.

The best known protocols in the TCP/IP suite are transmission control protocol (TCP), internet protocol (IP), file transfer protocol (FTP), simple mail transfer protocol (SMTP), and TELNET. The latter three are application layer protocols that correspond to FTAM, X.400, and VT in the OSI structure. TELNET is a protocol that allows users to log on a remote computer over the network and operate as if they were directly attached.

The user datagram protocol (UDP) is a connectionless version of TCP. It is used by simple network management protocol (SNMP), trivial file transfer protocol (TFTP), and versatile message transfer protocol (VMTP). See Chapter 38 for a discussion of SNMP. In addition, the protocol family includes address resolution protocol (ARP) and reverse address resolution protocol (RARP), which are discussed later. Routing information protocol (RIP) is used by UNIX-based computers for exchanging routing information, although *it* is